

FINAL
1N-34-CR
021 356

Visualization of
Unsteady Computational Fluid Dynamics

Final Technical Report
for
Grant # NAG2-884

Submitted

by

Robert Haimes

Computational Aerospace Sciences Laboratory
Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, MA 02139

March 1997

Introduction

The current compute environment that most researchers are using for the calculation of 3D unsteady Computational Fluid Dynamic (CFD) results is a super-computer class machine. The Massively Parallel Processors (MPPs) such as the 160 node IBM SP2 at NAS and clusters of workstations acting as a single MPP (like NAS's SGI Power-Challenge array and the J90 cluster) provide the required computation bandwidth for CFD calculations of transient problems.

If we follow the traditional computational analysis steps for CFD (and we wish to construct an interactive visualizer) we need to be aware of the following:

- **Disk space requirements**
A single snap-shot must contain at least the values (primitive variables) stored at the appropriate locations within the mesh. For most simple 3D Euler solvers that means 5 floating point words. Navier-Stokes solutions with turbulence models may contain 7 state-variables. The number can increase with the modeling of multi-phase flows, chemistry and/or electro-magnetic systems. If we examine a 5 equation system with 1 million nodes (with the field variables stored at the nodes) a single snap-shot will require 20 Megabytes. If 1000 time-steps are needed for the simulation (and the grid is not moving), 20 Gigabytes are required to record the entire simulation. This means that the workstation performing the visualization of this simulation requires vast amounts of disk space.
- **Disk speed vs. Computational speeds**
The time required to read the complete solution of a saved time frame from disk is now longer than the compute time for a set number of iterations from an explicit solver. Depending on the hardware and solver an iteration of an implicit code may also take less time than reading the solution from disk. If one examines the performance improvements in the last decade or two, it is easy to see that depending on disk performance (vs. CPU improvement) may not be the best method for enhancing interactivity. Workstation performance continues to double every 18 months. The performance of commodity drives has gone from about 1 Megabyte/sec in 1986 to about 5 Megabytes/sec in 1996.
- **Cluster and Parallel Machine I/O problems**
Disk access time is much worse within current parallel machines and cluster of workstations that are acting in concert to solve a single problem. In this case we are not trying to read the volume of data, but are running the solver and the solver outputs the solution. I/O is the bottleneck for a parallel machine with a front-end. The machine probably has the ability to compute in the GigaFLOP range but all this data has to be funneled to a single machine and put on disk by that machine. Clusters of workstations usually depend upon distributed file systems. In this case the disk access time is usually not the bottleneck, but the network becomes the pacing hardware. An IBM SP2 is a prime example of the difficulties of writing the solution out every iteration. The machine has a high-speed interconnect, but it is not currently used by the distributed file system. There are other access points into each node. Most SP2s have an Ethernet port for every node, some also have FDDI connections. These traditional network interfaces must be used for the file system.
- **Numerics of particle traces**
Most visualization tools can work upon a single snap shot of the data but some visualization tools for transient problems require dealing with time. One such tool is the integration of

particle paths through a changing vector field. After a careful numerical stability and accuracy analysis of integration schemes (funded by previous NAS contracts) it has been shown that there exist certain time-step limitations to insure that the path calculated is correct. Even for higher order integration methods, the limitation is on the order of the time step used for the CFD calculation. This is because of a physical limit, the time-scale of the flow. What this means (for the visualization system) is that in order to get accurate particle traces, the velocity field must be examined close to every time step the solver takes.

Because of the disk space requirements and the time to write the solution to disk, the authors of unsteady flow solvers perform some sort of sub-sampling. This sub-sampling can either be spatial or temporal. Because the traditional approach is to deal with the data as if it were many steady-state solutions, this sub-sampling I/O is almost always temporal. The individual running the simulation figures the frequency to write the complete solution based on the available disk space. In many cases, important transitions are missed. Also since the solution is coarsely sampled in time, streaklines (unsteady particle paths as discussed above) almost always produces erroneous results. The problem with sub-sampling is that the time-step selected for the visualization becomes based on the available disk space and not the physical problem.

pV3 Status

Work is in progress on a set of software tools designed specifically to address visualizing 3D unsteady CFD results in these super-computer-like environments. The above issues are resolved by co-processing the visualization. The visualization is concurrently executed with the CFD solver. The parallel version of **Visual3**, **pV3** required splitting up the unsteady visualization task to allow execution across a network of workstation(s) and compute servers. In this computing model, the network is almost always the bottleneck so much of the effort involved techniques to reduce the size of the data transferred between machines.

The following design goals for **pV3** have been met:

- **High Performance**
Take advantage of the proper hardware to get the best performance out of the entire compute arena. **pV3** requires graphics hardware so that scene rendering time is not a limitation and the data presented to the investigator is of high quality and timely. Also, most visualization techniques are embarrassingly parallel (based on elements within the computational volume). The execution of these tools is done within the partitioning performed to parallelize the CFD solver.
- **Interactive**
The goal of any scientific visualization package should be to allow the assimilation of the vast amounts of data produced by the models and solvers in order to better understand the underlying physics. The ultimate goal, with this new knowledge, is to affect design and produce a better car, aircraft, gas-turbine engine, etc. This can only be done by interactively poking and probing into the data to interrogate areas of interest.
- **Co-processing**
An important part of **pV3** is the ability to visualize the data as the solver or model progresses in time. It is also designed to allow the solver to run as independently as possible. If the solution procedure takes hours to days, **pV3** can *plug-into* the calculation, allow viewing of the data as it changes, then can *unplug* with the worst side-effect being the temporary allocation of memory and a possible load imbalance.

- **Visual3** functionality and programming
pV3 provides the same kind of functionality as **Visual3** with the same suite of tools and probes. The data represented to the investigator (the 3D, 2D and 1D windows with cursor mapping) is the same. Also the same Graphical User Interface (GUI) is used.

For the desired flexibility and the merging of the visualization with the solver, some programming is required. The coding is simple; like **Visual3**, all that is required of the programmer is the knowledge of the data. Learning the details of the underlying graphics, data extraction, and movement (for the visualization) is not needed. If the data is distributed in a cluster of machines, **pV3** deals with this, resulting in few complications to the user.

pV3 Rev 1.20 was released in October 1996. It is anticipated that Rev 1.25 will be released sometime shortly after the end of this contract. This port will include support for multiple interactive viewers and concurrent batch subsystem execution within a complex mix of solvers. **pV3** Rev 1.25 will also include the support of the work described in the Status Section. The following machines are (and will be) supported as 'clients' (the computers containing the volume of data and performing the solver):

- CRAY J90s and C90s
- DEC Alphas running DEC Unix
- HP 9000/700 series at HP-UX 9.0 (or higher)
- IBM RS/6000s including the SP2s
- SGIs IRIX 5.x and 6.x in 32 bit mode
- SGIs (R8000s and R10000s) running IRIX 6.x in 64 bit mode
- SUNs running Solaris

pV3 supports interactive viewers and batch servers for the following workstations (that have 3D graphics hardware) and support OpenGL:

- DEC Alphas running DEC Unix
- IBM RS/6000s
- SGIs IRIX 5.x and 6.x in 32 bit mode
- SGIs (R8000s and R10000s) running IRIX 6.x in 64 bit mode
- SUNs running Solaris

Note: HPs are not supported because of their lack of OpenGL and multi-threading support.

Presentations

A Modular Approach to Visualization for Parallel CFD Applications.
 Parallel CFD '96, Capri Italy, May 1996.

Real-Time Visualization of an HPF-based CFD Simulation (with M. Kremenetsky and A. Vaziri).
 Parallel CFD '96, Capri Italy, May 1996.

Visualization in a Parallel Processing Environment (with D. Edwards) -- invited.
 AIAA Aerospace Sciences Meeting & Exhibit, January 1997. AIAA Paper 97-0348.

Demonstrations

Supercomputing '96, Pittsburgh, November 1995.

AIAA Aerospace Sciences Meeting & Exhibit, Reno, January 1997.

Status of this Work

pV3 is currently being used for parallel CFD solver development and debugging as well as for quality assessment for long running applications. There is promise that this technology will also be important in production environments with the advent of the 'batch' subsystem. **pV3**'s development is progressing in a modular fashion where groups of interconnected tasks can be selected based on the underlying CFD computation, on what equipment it is executing upon, and the type of execution (interactive, batch or both).

The visualization efforts at NAS has been re-directed, also towards a building-block approach to visualization applications (where **elVis** is the first example). This modular programming will provide custom visualization solutions with a minimum of development time.

pV3's co-processing capabilities are unique within the CFD visualization community. There have been years of development and testing (as well as user feedback) in an environment closely coupled to solver technologies. The work done under this contract was the development of an Application Programming Interface (API) to the co-processing modules of **pV3** so that other visualization modules (or complete systems) can have direct access to the running solver. The goal of the work was to include portions of **pV3** as part of NASA's library of visualization building-blocks. The API is documented in the appended manual: "Server Builders Guide".

This work progressed in the following stages:

- Interface Specification, Design and Documentation
The API was specified, documented and then reviewed by the visualization group at NASA Ames. A meeting was held early in this work that brought together MIT and NAS personnel for the preliminary design.
- Interface Development
pV3 was modified to support the API.
- Interface Test
A second trip was used to couple **FEL** (the Field Encapsulation Library) with **pV3**. This successfully displayed that **pV3** client-side could co-exist with NAS' object-oriented library and could be included with a solver. **pV3**'s viewer was used to display the results.
- Software Release
The API will become part of the next **pV3** distribution. The document mentioned above will be included, so that other developers can take advantage of the **pV3** co-processing modules.

Finally, this approach allows for what was intended in an earlier phase of this contract (trying to use **fGL** for the disk based visualization extracts). The MIT and NAS efforts have been closely coupled. Users will now be able to choose from either the **pV3** viewers or the easier-to-use products from the NAS visualization team. And either **pV3**'s data extractors can be used or code like **FEL** or both at the data/solver end.

Server Builder's Guide

for
pV3 Rev. 1.25

Bob Haines

March 19, 1997

License

This software is being provided to you, the LICENSEE, by the Massachusetts Institute of Technology (M.I.T.) under the following license. By obtaining, using and/or copying this software, you agree that you have read, understood, and will comply with these terms and conditions:

Permission to use, copy, modify and distribute, this software and its documentation for any purpose and without fee or royalty is hereby granted, provided that you agree to comply with the following copyright notice and statements, including the disclaimer, and that the same appear on ALL copies of the software and documentation:

Copyright 1996-1997 by the Massachusetts Institute of Technology. All rights reserved.

THIS SOFTWARE IS PROVIDED "AS IS", AND M.I.T. MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED. BY WAY OF EXAMPLE, BUT NOT LIMITATION, M.I.T. MAKES NO REPRESENTATIONS OR WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE LICENSED SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

The name of the Massachusetts Institute of Technology or M.I.T. may NOT be used in advertising or publicity pertaining to distribution of the software. Title to copyright in this software and any associated documentation shall at all times remain with M.I.T., and USER agrees to preserve same.

Contents

1	Introduction	5
1.1	Unsteady Classification	5
1.2	Extracts	5
1.3	Multi-Threading	7
1.4	Programming Notation	8
2	Built-in Server Side Extracts	9
2.1	Surfaces	9
2.2	StreamLines	12
2.3	Particles	14
2.4	Vector Clouds	16
3	Calls that can be Invoked by Either Thread	17
3.1	pV_DiscipStat	17
3.2	pV_ClientStat	17
3.3	pV_FieldStat	18
3.4	pV_CutStat	18
3.5	pV_SetDiscip	19
3.6	pV_GetExtract	19
4	I/O Thread Routines and Calls	22
4.1	pVSafe	22
4.2	pV_Register	22
4.3	pVSetExtract	24
4.4	pV_SetExState	24
4.5	pV_CrExtract	25
4.6	pV_DeExtract	26
5	Graphics Thread Routines and Calls	27
5.1	pVRender	27
5.2	pV_GetSub	27

6	Running Your Server	28
6.1	Environment Variables	28
6.2	Special File - The Lock File	28
7	An FEL Example	29
7.1	Server code	29
7.1.1	I/O Thread code	29
7.1.2	Graphics Thread code	31
7.2	Client Side	32
A	Plotting Masks for Built-in Extracts	35
A.1	Cut Surfaces	35
A.2	StreamLines	35
A.3	Particles	36
A.4	Vector Clouds	36

1 Introduction

This manual is a guide for those individuals wishing to use **pV3**'s client-side API and network based data movement, but do not want to view the data via **pV3**'s interactive server. This may be necessary when either the data presentation is not appropriate or some other workstation-enhanced technique (such as a Cave or VR) is used.

1.1 Unsteady Classification

Transient applications are classified in the following manner:

- $OPT = 0$; *Steady-State*

This is the simplest case. Nothing changes in time.

- $OPT = 1$; *Data Unsteady*

In this type of application the grid structure and position are fixed in time. The data defined at the nodes (both scalar and vector) changes with each time step. An example is when a boundary condition in the domain is changing.

- $OPT = 2$; *Grid Unsteady*

These cases are 'Data Unsteady' plus the grid coordinates associated with each node are also allowed to move with each snapshot. An example of this is stator/rotor interaction in turbomachinery. The stator and rotor grids are separate, with the rotor grid sliding past the stator grid. In this case the stator portion is actually 'Data Unsteady' and the rotor grid moves radially.

- $OPT = 3$; *Structure Unsteady*

If the number of nodes, number of cells or cell connectivity changes from iteration to iteration the case is 'Structure Unsteady'. An example of this mode is store separation.

1.2 Extracts

pV3 has been designed to minimize network traffic. The client-side library extracts lower dimensional data required by the requested visualization tool from the volume of data in place. This distilled data is transferred to the graphics workstation. To further reduce the communication burden posed by the visualization session, the transient problem classification described above is used. Only the extracted data that has changed from the last iteration is sent over the network. An extract is therefore the results of a visualization tool (geometric cut, iso-surface, streamline and etc.) collected in a manner that provides

flexibility and minimizes the volume of data so that a network based visualization system can provide good frame-rates.

There are 2 type of extracts allowed within the **pV3** system. The first are pre-defined and include those listed in Section 2. At startup the following built-in extracts will have been created:

- Domain Surface. One extract for each global surface with the plotting attributes set by the client side.
- Particle Data. For unsteady flows with vector fields.
- Vector Cloud. For clients with vector fields. The attributes set for this tool are in the off state at initialization.
- Dynamic. One extract is used by *pV3Server* for all dynamic tools.

The second type of extract is programmer defined. In this case code must be supplied at both the client and server sides.

Each extract is divided into as many as 12 sub-extract components. This is done for the following reasons:

- Unsteady updates
Extract information should be segregated based upon reducing the amount of data. For example, a geometric-based cut is generated in a *Data Unsteady* client. For the second and subsequent iterations, all geometric data does not change, and therefore need not be retransmitted. The only new data required is the scalar field values associated with each vertex of the object. With only moving this data, the extract can be properly rendered, coloring the surface based on the current field. In this case, having the scalar values as a sub-extract produces a great benefit.
- Network usage
Segregating data based on type is important. The underlying message passing (i.e. PVM) may deal with hetrogenous machine environments. Being able to do the bit or byte *twiddling* – based on type is required.

See Section 2 for an example of **pV3**'s pre-defined extracts and their sub-extract types.

For programmer defined extracts that use derived types and/or complex structures, the data must be decomposed at the client side so that it can pass through the network interface. On the server side the data can be reassembled from the sub-extracts.

1.3 Multi-Threading

pV3's servers (*pV3Server*, *pV3Batch* and *pV3Viewer*) are all multi-threaded. In fact, they have 2 threads. The application that gets built using this guide also uses 2 threads, like the interactive server. See Figure 1.




Figure 1: The **pV3** Server's Threading Control

All extracts are double-buffered. This allows concurrent execution of the threads without data contention. The I/O thread collects extracts in the client's current iteration as the Graphics thread is rendering from the previous set of data. Once the thread handshaking is complete, the buffers are swapped, the last rendered data is thrown away (where appropriate) and the process continues.

The application built from this guide will have the same architecture as seen in Figure 1. The difference is that the **pV3** code used to 'Render scene', 'Get Xevents' and 'Update state' is replaced by a single programmer supplied routine; **pVRender**. See Section 5.1.

1.4 Programming Notation

pV3 was designed to be accessible from both FORTRAN and C. FORTRAN is more restrictive in argument passing and naming, therefore it has shaped the programming interface. The routine descriptions in this guide are from the C programmer's point of view. But because FORTRAN is supported with the same API all routine arguments are pass by reference. It is assumed that a routine's argument is not modified unless documented as such.

For IBM and HP ports, all **pV3** entry points are the FORTRAN names in lower-case. On all other platforms except the CRAY, external entries are lower-case with an underscore ('_') appended to the end. CRAY entry points are upper-case with no appended underscores. See the file 'pV3ser.h' or 'wsdepend.h' in the servers subdirectory of the distribution for a method to avoid these problems.

Consistant with the **pV3** naming convension, the routines that are part of **pV3**'s server suite are prefixed with 'pV_', those that are supplied by the programmer start with 'pV'.

2 Built-in Server Side Extracts

The following section describes the internal data stored in the **pV3** server structures for the built-in extracts. This data can be used to produce the graphics objects that get rendered to make the scene. Each tool generates a different type of *extract* from the 3D data in the client(s). The data gets transmitted to the server and is stored for as long as it is needed. Each *extract* consists of a number of *sub-extract* types, and there is a complete collection of *sub-extracts* for each client. Note: each client's data is stored separately.

2.1 Surfaces

This data is generated by the **pV3** scalar tools (planar cuts, programmed cut surfaces, iso-surfaces and domain surfaces). This data is exposed so that new 'probes' may be easily generated. The size of many of these arrays (and therefore the pointers) will change during the execution of **pV3**, so when using this data, get the current pointers before accessing the memory.

Extract	Type	Valid Sub-Extracts
2	Planar Cut	0 1 2 3 4 5 6 7
4	Geometric Cut	0 1 2 3 4 5 6 7 8
5	Domain Surface	0 1 2 3 4 5 6 7 8
7	Iso-Surface	0 1 2 3 4 5 6 7

0 - Surface Sub-Extract *Tris*

The following data defines the disjoint triangle space. Where the number of triangles in the structure is KTRI.

int TRIS[KTRI][4]

disjoint triangle definitions.

TRIS[][0] = first node index for the triangle.

TRIS[][1] = second node index for the triangle.

TRIS[][2] = third node index for the triangle.

TRIS[][3] = the parent 3D cell number (in the client).

1 - Surface Sub-Extract *Quads*

The following data defines the disjoint quadrilateral space. Where the number of quadrilaterals in the structure is KQUAD.

int QUADS[KQUAD][5] disjoint quadrilateral definitions.

QUADS[][0] = first node index for the quadrilateral.
QUADS[][1] = second node index for the quadrilateral.
QUADS[][2] = third node index for the quadrilateral.
QUADS[][3] = fourth node index for the quadrilateral.
QUADS[][4] = the parent 3D cell number (in the client).

2 - Surface Sub-Extract *XYZ*

The following data defines the 3D coordinates for the nodes (and therefore also the number of nodes) that support the surface. The number of nodes in the structure is KXYZ.

float XYZ[KXYZ][3] (x, y, z)-coordinates for the nodes.

3 - Surface Sub-Extract *Mesh*

The following data defines the disjoint lines that make-up the intersection of the cell edges and the cutting surface. The number of line segments in the structure is KFACE.

int FACE[KFACE][2] disjoint line definitions.

FACE[][0] = first node index for the line.
FACE[][1] = second node index for the line.

4 - Surface Sub-Extract *Outline*

The following data defines the disjoint lines that make-up the outline of the surface. The number of line segments in the structure is KEDGE.

int EDGE[KEDGE][3] disjoint line definitions.

EDGE[][0] = first node index for the line.
EDGE[][1] = second node index for the line.
EDGE[][2] = the parent surface face number (in the client).

5 - Surface Sub-Extract *Scalar*

The following data defines the current scalar for the nodes (and therefore also the number of nodes) that support the surface. The number of nodes in the structure is KS and is the same as KXYZ.

float S[KS] scalar functional values for the nodes.

6 - Surface Sub-Extract *Vector*

The following data defines the current vector for the nodes (and therefore also the number of nodes) that support the surface. The number of nodes in the structure is KV and is the same as KXYZ.

float V[KV][3] vector values (V_x, V_y, V_z) for the nodes.

7 - Surface Sub-Extract *Threshold*

The following data defines the current threshold values for the nodes that support the surface. The number of nodes in the structure is KT and is the same as KXYZ.

float T[KT] threshold functional values for the nodes.

8 - Surface Sub-Extract *2D Mapping*

The following data defines the 2D mapping for the nodes that support the surface. The number of nodes in the structure is KXY and is the same as KXYZ.

float XY[KXY][2] raw (x', y')-coordinates as specified by the client.

Notes:

- (1) The 2D mapping for planar cuts is implicit and not required from the client.
- (2) There is no 2D mapping for iso-surfaces.

2.2 StreamLines

This data is generated by the **pV3** clients during the integration of instantaneous streamlines. The size of many of these arrays (and therefore the pointers) will change during the execution of **pV3**, so when using this data, get the current pointers before accessing the memory. Unlike all other extracts, the number of sub-extracts is not a function of the number of clients but of the maximum allotted streamline segments (that is greater than the number of clients). This allows a streamline to reenter a client more than once.

0 - StreamLine Sub-Extract *Cell*

The following data contains the 3D cell number for the position of the point for this segment (used for the point probe). The number of entries in the structure is KCELL and is the same as KXYZ.

int CELL[KCELL] the parent 3D cell number (in the client).

1 - StreamLine Sub-Extract *Time*

The following data defines the integration pseudo-time for the point (used for streamline animation). Where the number of elements in the structure is KTIME and is the same as KXYZ.

float TIME[KTIME] integration time (from the seed position).

2 - StreamLine Sub-Extract *XYZ*

The following data defines the 3D coordinates for the points that support this poly-line segment. The number of nodes in the structure is KXYZ.

float XYZ[KXYZ][3] (x, y, z)-coordinates for the points.

3 - StreamLine Sub-Extract *Div*

The following data defines the cross-flow divergence felt by each point during the integration. Where the number of elements in the structure is KDIV and this is the same as KXYZ.

float DIV[KDIV] used for streamtube rendering, where the size of the tube is based on a starting size multiplied by e to this power.

4 - StreamLine Sub-Extract *Angle*

The following data contains the curl for each point, calculated during the integration, in this segment of the streamline Where the number of entries in the structure is KANG and this is the same value as KXYZ.

float ANG[KANG] angle of the twist for ribbons in degrees.

5 - StreamLine Sub-Extract *Scalar*

The following data defines the current scalar for the points that support the line in this segment. The number of points in the structure is KS and this is the same as KXYZ.

float S[KS] scalar functional values for the points.

6 - StreamLine Sub-Extract *Vector*

The following data defines the current vector for the points that make up this segment of the streamline. The number of elements in the structure is KV and this is the same as KXYZ.

float V[KV][3] vector values (V_x, V_y, V_z) for the points.

7 - StreamLine Sub-Extract *Threshold*

The following data defines the current threshold values for the points that support the poly-line. The number of entries in the structure is KT and is the same as KXYZ.

float T[KT] threshold functional values for the points.

2.3 Particles

This data is updated by the **pV3** clients during the bubble integration at each time-step. The size of many of these arrays (and therefore the pointers) will change during the execution of **pV3**, so when using this data, get the current pointers before accessing the memory.

0 - Particle Sub-Extract *Number*

The following data contains the unique particle number for each bubble in that client. The number of entries in the structure is **KNUM** and this is the same as **KXYZ**.

int NUM[KNUM] the global particle number.

1 - Particle Sub-Extract *Time*

The following data defines the start time for each bubble. The number of elements in the structure is **KTIME** and this number is the same as **KXYZ**.

float TIME[KTIME] bubble simulation time when the particle was seeded.

2 - Particle Sub-Extract *XYZ*

The following data defines the current 3D coordinates for the particles. The number of nodes in the structure is **KXYZ**.

float XYZ[KXYZ][3] (x, y, z)-coordinates for the bubbles.

3 - Particle Sub-Extract *Div*

The following data defines the cross-flow divergence currently felt by each bubble. Where the number of elements in the structure is **KDIV** and this is the same as **KXYZ**.

float DIV[KDIV] optionally used for bubble rendering, where the size of the particle is based on a starting size multiplied by **e** to this power.

5 - Particle Sub-Extract *Scalar*

The following data defines the current scalar for the particles in this client. The number of points in the structure is KS and this is the same as KXYZ.

float S[KS] scalar functional values for the bubbles.

6 - Particle Sub-Extract *Vector*

The following data defines the current vector for the particles. The number of elements in the structure is KV and this number is the same as KXYZ.

float V[KV][3] vector values (V_x, V_y, V_z) for the bubbles.

7 - Particle Sub-Extract *Threshold*

The following data defines the current threshold values for the particles. The number of entries in the structure is KT and is the same as KXYZ (the number of bubbles).

float T[KT] threshold functional values for the bubbles.

10 - Particle Sub-Extract *Group Index*

The following data defines the current group number for the particles. The number of entries in the structure is KGI and is the same as KXYZ (the number of bubbles).

int GI[KGI] group index for the bubbles (used for *time lines*).

2.4 Vector Clouds

2 - VC Sub-Extract *XYZ*

The following data defines the coordinates for the 3D nodes that satisfy the threshold limits within each client. The number of nodes in the structure is KXYZ.

float XYZ[KXYZ][3] (x, y, z) -coordinates for the vector cloud.

5 - VC Sub-Extract *Scalar*

The following data defines the current scalar for the vector cloud. The number of points in the structure is KS and this number is the same as KXYZ.

float S[KS] scalar functional values for the 3D nodes.

6 - VC Sub-Extract *Vector*

The following data defines the current vector for each node in the client that satisfies the threshold limits. The number of elements in the structure is KV and this number is the same as KXYZ.

float V[KV][3] vector values (Vx, Vy, Vz) for the vector cloud.

3 Calls that can be Invoked by Either Thread

3.1 pV_DiscipStat

PV_DISCIPSTAT(DID,NID,NCL,DNAME)

This routine returns the status for the disciplines known to the system.

int *DID	The discipline ID. The first discipline id always 0, therefore it is always safe to make this call with this argument zero.
int *NID	The total number of disciplines in the simulation. Returned.
int *NCL	The number of clients in the discipline DID. Returned.
char DNAME[20]	The discipline's name. Returned.

NID must be atleast 1 for any valid **pV3** application therefore it is always valid to make this call with DID = 0.

3.2 pV_ClientStat

PV_CLIENTSTAT(DID,CID,OPT,NCL,ON,CNAME)

This routine returns the status for a client within a discipline.

int *DID	The discipline ID. Must be a value from 0 to NID-1.
int *CID	The client ID. Must be a value from 1 to NCL.
int *OPT	The client's unsteady index (0-3). Returned.
int *NCL	The number of clients in the discipline DID. Returned.
int *ON	Visibility flag -- not used. Returned.
char CNAME[20]	The clients's name. Returned.

NCL must be atleast 1 for any valid **pV3** discipline therefore it is always valid to make this call with CID = 1.

3.3 pV_FieldStat

PV_FIELDSTAT(DID,FID,NFL,FTY,FLIMS,FNAME)

This routine returns the status for a field variable within a discipline.

int *DID	The discipline ID. Must be a value from 0 to NID-1.
int *FID	The field ID. Must be a value from 1 to NFL.
int *NFL	The number of field variables in the discipline DID. Returned. An error is indicated by the value -1 which indicates that the discipline index or field ID is out of range.
int *FTY	The field's type (1-5). Returned. 1 Scalar 2 Vector 3 Surface scalar 4 Surface vector 5 Threshold
float FLIMS[2]	Field limits from the clients. Returned.
char FNAME[32]	The field's name. Returned.

NFL must be atleast 1 for any valid **pV3** application therefore it is always valid to make this call with FID = 1.

3.4 pV_CutStat

PV_CUTSTAT(DID,CIN,NCT,CTITLE)

This routine returns the status for a programmed cut within a discipline.

int *DID	The discipline ID. Must be a value from 0 to NID-1.
int *CIN	The cut index. Must be a value from 1 to NCT.
int *NCT	The number of geometric cuts in the discipline DID. Returned.
char CTITLE[32]	The geometric cut's title. Returned.

There may be no cuts (i.e. $NCT = 0$). To determine the number of cuts, call this routine with CIN = 1, then check NCT.

3.5 pV_SetDiscip

PV_SETDISCIP(DID)

This routine sets the current discipline. Not needed for a single discipline case.

int *DID	The discipline ID. Must be a value from 0 to NID-1.
----------	---

3.6 pV_GetExtract

PV_GETEXTRACT(EX,TYPE,EXNUM,IVEC,RVEC,NAME,NEXTEX)

Returns the internal **pV3** extract structure info for the current discipline. The extracts form a linked list. There is a unique list for each discipline. This routine allows the scanning of all active extracts by continual calls until the desired extract is found.

void **EX	Extract pointer. On input, this is the desired extract. The special case of the first extract is indicated by a NULL and is updated with the actual extract pointer.
int *TYPE	The extract type. Returned.
int *EXNUM	The extract number. Returned.
int IVEC[]	Integer data set based on TYPE (length also determined by TYPE). Returned.
int RVEC[]	Float data set based on TYPE (length also determined by TYPE). Returned.
char NAME[20]	Extract name. Returned.
void **NEXTEX	Returned pointer to the next extract. NULL indicates that this is the last extract. NEXTEX can be used in the next call to pV_GetExtract (argument EX) to continue scanning the list.

The following data is related to data in the Graphics buffer:

- Planar Cut - TYPE = 2

IVEC[0] = Plot Mask

IVEC[1] = Scalar field index

IVEC[2] = Vector field index

IVEC[3] = Threshold index

RVEC[0-8] = Cut corners - Three of the 4 corners that denote the plane

RVEC[9-11] = Plane normal

- Geometric Cut - TYPE = 4

IVEC[0] = Plot Mask

IVEC[1] = Scalar field index

IVEC[2] = Vector field index

IVEC[3] = Threshold index

IVEC[4] = Cut index

RVEC[0] = Z prime

- Domain Surface - TYPE = 5

IVEC[0] = Plot Mask

IVEC[1] = Scalar field index

IVEC[2] = Vector field index

IVEC[3] = Threshold index

IVEC[4] = Mapping flag

IVEC[5] = Special surface scalar index

IVEC[6] = Special surface vector index

- Iso-Surface - TYPE = 7

IVEC[0] = Plot Mask

IVEC[1] = Scalar field index

IVEC[2] = Vector field index

IVEC[3] = Threshold index

IVEC[4] = Scalar index for Iso-Surface

RVEC[0] = Z prime

- StreamLine - TYPE = 18

IVEC[0] = Plot Mask
 IVEC[1] = Scalar field index
 IVEC[2] = Vector field index
 IVEC[3] = Threshold index
 IVEC[4] = StreamLine Group number
 IVEC[5] = Client-id for client with seed location
 IVEC[6] = Cell index in client to start StreamLine
 IVEC[7] = Minimum StreamLine number for group
 IVEC[8] = Maximum StreamLine number for group
 IVEC[9] = Surface Index (0 - volume StreamLine)
 IVEC[10] = Number of StreamLine segments
 RVEC[0-2] = Seed location (XYZ)

- Particles - TYPE = 19

IVEC[0] = Plot Mask
 IVEC[1] = Scalar field index
 IVEC[2] = Vector field index
 IVEC[3] = Threshold index

- Vector Cloud - TYPE = 20

IVEC[0] = Plot Mask
 IVEC[1] = Scalar field index
 IVEC[2] = Vector field index
 IVEC[3] = Threshold index
 RVEC[0] = Threshold minimum
 RVEC[1] = Threshold maximum

- Programmer-defined - TYPE > 100

IVEC[0] = Plot Mask - Not used
 IVEC[1] = Scalar field index
 IVEC[2] = Vector field index
 IVEC[3] = Threshold index
 IVEC[4] = IVAL
 RVEC[0-8] - Float values associated with the extract

4 I/O Thread Routines and Calls

4.1 pVSafe

PVSAFE()

This programmer-supplied routine is called when the graphics thread of the server is stalled. This is the time where calls can be made that require neither thread to be active. The buffers have not been swapped, so that queries of extracts will look at the last state.

No Arguments

NOTE:

The first call to **pVSafe** is done when there is only one thread. This should be used to register all extracts (at least for the first allocation).

4.2 pV_Register

EX = PV_REGISTER(INDEX,NAME,SUBTYPE,SUBSIZE,SUBOPT, SUBLOC,ROUTINE,EXNUM)

Registers a programmer-defined extract with the **pV3** server. This routine should only be called when the threads are sync'ed, therefore the only valid place to execute this routine is within **pVSafe**. For multi-disciplinary cases, the discipline index must be set so that the extract is registered within the appropriate discipline. The client-side extraction code must be linked with the client application. See the Advanced Programmer's Guide.

void *EX	The extract pointer if EXNUM does not indicate an error.
int *INDEX	The extract index. This number must be greater than 100 and defines an extract. Different disciplines must use unique extract indices!
char NAME[20]	Extract name.
int SUBTYPE[12]	The subextract types. Each extract is composed of up to 12 subextracts for each client. This vector defines whether the subextract is an integer (0) or a float (1).
int SUBSIZE[12]	The subextract size per length. For example, if the subextract is for the 3D coordinates (X,Y,Z) that support the extract, the size would be 3.

int SUBOPT[12]

The level of unsteadyness that requires the data at every time-step. Valid entries are 0 to 2. The following table specifies what action is taken with an existing subextract:

Client's OPT - >	0	1	2	3
SUBOPT = 0	leave	refill	refill	refill
SUBOPT = 1	leave	leave	refill	refill
SUBOPT = 2	leave	leave	leave	refill

int SUBLOC[12]

The subextract's locality. If this subextract comes from the clients then the value is -1. If this subextract is local to the server and it's length is set by another subextract, then SUBLOC must contain the index (0 biased) to that extract. SUBOPTs for local subextracts must match that of the keyed subextract.

void (*ROUTINE)()

Not used – for compatibility with the Advanced Programmer's interface.

int *EXNUM

This is a status return. If the value is zero or greater, that indicates success. The value is the number used for multiple allocations of extracts with the same INDEX. If the number is negative it is an indication of an error:

- 1 - Invalid INDEX number
- 2 - Invalid SUBTYPE in one of the entries
- 3 - Invalid SUBSIZE in one of the entries
- 4 - Invalid SUBOPT in one of the entries
- 5 - Invalid SUBLOC in one of the entries
- 6 - SUBTYPE mismatch for subsequent calls using INDEX
- 7 - SUBSIZE mismatch for subsequent calls
- 8 - SUBOPT mismatch for subsequent calls
- 9 - SUBLOC mismatch for subsequent calls
- 10 - ROUTINE mismatch for subsequent calls
- 11 - Allocation error
- 12 - Routine not called from pVSafe
- 13 - SUBOPTs mismatch for local subextract

4.3 pVSetExtract

PVSETEXTRACT(INDEX,EXNUM,PLOTMASK,REQMASK,IVAL,RVEC)

This routine gets called for each registered extract during the request collection phase.

int *INDEX	The extract index. This number must be greater than 100 and defines an extract (and discipline).
int *EXNUM	The extract number associated with INDEX.
int *PLOTMASK	Not used, for compatibility with the Advanced Programmer's interface.
int *REQMASK	The request mask. Each bit specifies which subextracts are required to satisfy the plotting attributes. For example, 5 requests subextract 0 and subextract 2. If the most-significant bit is set all subextracts are requested, even if based on SUBOPT, the data exists (i.e. some state has changed). Must be filled on return.
int *IVAL	An integer sent to the clients associated with this extract. Must be set upon return.
float RVEC[9]	A float vector of data sent to the clients with the request for this extract. Must be set upon return.

4.4 pV_SetExState

PV_SETEXSTATE(EX,PLTMASK,SCALAR,VECTOR,THRES)

Sets the plot mask and field variables for the specified extract. This routine should be called from **pVSafe** to insure that the change to the attributes is effective for the next iteration.

void **EX	Extract pointer as returned by pV_GetExtract , pV_Register or pV_CrExtract . NOTE: these pointers do not change during the life of the server application.
int *PLTMASK	Plot mask to be set for the extract – built-ins only. See the Appendix for the mask values.
int *SCALAR	Scalar field index to be used.
int *VECTOR	Vector field index to be used.
int *THRESH	Threshold index, (-) indicates a scalar field index.

4.5 pV_CrExtract

EX = PV_CREXTRACT(TYPE,NAME,IVEC,RVEC,EXNUM)

Creates a built-in extract for the current discipline.

void *EX	The extract pointer if EXNUM does not indicate an error.
int *TYPE	The extract type. Valid types are 2, 4, 7 and 18.
char NAME[20]	Extract name.
int IVEC[6]	Integer data set based on TYPE.
float RVEC[9]	Float data set based on TYPE.
int *EXNUM	The returned extract instance. If the number is negative it is an indication of an error: -2 - Invalid TYPE -11 - Allocation error -12 - Routine not called from pVSafe

- Planar Cut - TYPE = 2

RVEC[0-8] = Cut corners - Three of the 4 corners that denote the plane

- Geometric Cut - TYPE = 4

IVEC[0] = Cut index - 1 to NCT

IVEC[1] = Instancing Mask - 0 allow replication, 1 no instancing

RVEC[0] = Z prime

- Iso-Surface - TYPE = 7

IVEC[0] = Scalar index for Iso-Surface

RVEC[0] = Z prime

- StreamLine - TYPE = 18

IVEC[0] = StreamLine Group number

IVEC[1] = Client-id for client with seed location

IVEC[2] = Cell index in client to start StreamLine

IVEC[3] = Minimum StreamLine number for group

IVEC[4] = Maximum StreamLine number for group

IVEC[5] = Surface Index (0 - volume StreamLine)

RVEC[0-2] = Seed location (XYZ)

4.6 pV_DeExtract

PV_DEEXTRACT(EX)

Deletes the specified extract. This routine should only be called from pVSafe.

void **EX Extract pointer as returned by pV_GetExtract,
 pV_Register or pV_CrExtract.

5 Graphics Thread Routines and Calls

5.1 pVRender

PVRENDER(TIME)

This is the routine that gets called to render the data. It might get called more than once for each set of data depending on the length of time to render and the client-side update frequency. In this case, the value of TIME does not change.

float *TIME	The simulation time for the data.
-------------	-----------------------------------

5.2 pV_GetSub

PV_GETSUB(EX,SUBEX,NUMCS,PTR,LEN,CID)

Returns the internal pV3 sub-extracts. This routine returns the Graphics thread pointers (from the two buffers).

void **EX	Extract pointer as returned by pV_GetExtract, pV_Register or pV_CrExtract.
int *SUBEX	Sub-extract number (0-11 based on TYPE).
int *NUMCS	Client index or StreamLine segment number (0 biased).
void **PTR	Returned pointer to the structure. NULL indicates that the memory block is not allocated.
int *LEN	Length of structure. A 0 (zero) indicates that the structure is not currently filled. Returned.
int *CID	Client-id for the client that produced the segment Returned (StreamLines Only).

6 Running Your Server

The **PVM** daemon(s) and with co-processing, the solver, must be executing. Without a **pV3** server running, every time the solution is updated, a check is made for the number of members in the **PVM** group *pV3Server* (Note: this name can be changed for multiple jobs running under the same user ID – see the Section 6.1 for the environment variable ‘pV3_Group’). If no servers are found, no action is taken. When a **pV3** server starts, it enrolls in the specified group. The next time the solution is updated, an initialization message is processed and the visualization session begins. Each subsequent time in the solver completes a time step, visualization state messages and *extract* requests are gathered, the appropriate data calculated, collected and sent to the active server(s).

When the user is finished with the visualization, the server sends a termination message and exits. The clients receive the message, and if no other servers are running, cleans up any memory allocations used for the visualization. Then the scheme reverts to looking for server initialization, if termination was not specified at **pV3** client initialization.

6.1 Environment Variables

A **pV3** server built using this guide automatically looks at two Unix environment variables:

‘pV3_TO’ should be used to change the internal Time-Out constant. If the variable is set, it must be an integer string which is the number of seconds to use for the Time-Out constant (the server’s default is 60). This may be required if the time between solution updates is long. See the section in the **pV3** Server User’s Reference Manual on *Time-Outs and Error Recovery*.

‘pV3_Group’ is usefull for differentiating multiple **PVM** jobs running under the same user ID. If this variable is set for the solver (client-side) before execution, it overrides the default client side group name *pV3Client*. The name used is the string assigned to this variable with *Client* appended. By setting this variable before server execution, it will set the server group to the variable’s string with *Server* appended instead of using *pV3Server*. Only clients with the appproaite matching group name will be connected to this session.

6.2 Special File - The Lock File

If the server is running on a multi-processor SGI workstation (PowerSeries, Onyx or PowerOnyx) a file is used for the coordination of the 2 threads generated during execution. This file has the name ‘.pV3.locks’ and is open in the current directory. It should be noted that running two invocations of the **pV3** server from the same directory will NOT work. Both will use the same file for the lock and semaphore arena!

7 An FEL Example

The following is a very simple coding example of both making a programmed defined extract (at the server and client-side) as well as skeleton code for the customized server. It is assumed that the visualization is steady-state and there is only one discipline.

7.1 Server code

7.1.1 I/O Thread code

```
#include <stdio.h>
#include <stdlib.h>
#include "pV3ser.h"

/* required to deal with visualization control and state */

extern void get_field(int *scalar, int *vector, int *thresh);
extern int  get_pltmask(int type, int exnum);

void
PVSAFE()
{
    static int EXNUM = -14;
    static char *name = "FEL StreamLine      ";
    static int subtype[12] = { 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0};
    static int subsize[12] = { 3, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1};
    static int subopt[12]  = { 0, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1};
    static int subloc[12]  = {-1,-1,-1,-1,-1,-1,-1,-1,-1,-1,-1,-1};

    int  ivec[11], type, exnum, pltmask, scalar, vector, thresh, opt;
    char  exname[20];
    float rvec[12];
    void  *ex, *nextex;

    /* register the extract */
    opt = 101;
    if (EXNUM != -14) PV_REGISTER(&opt, name, subtype, subsize,
                                subopt, subloc, NULL, &EXNUM);
}
```

```

/* get the current field variables */
get_field(&scalar, &vector, &thresh);

/* loop through all defined extracts */
ex = NULL;
PV_GETEXTRACT(&ex, &type, &exnum, ivec, rvec, exname, &nextex);
while (ex != NULL) {
    PV_GETEXTRACT(&ex, &type, &exnum, ivec, rvec, exname, &nextex);

    /* set extract's attributes for the next set of data */
    pltmask = 0;
    if (type < 100) pltmask = get_pltmask(type, exnum);
    PV_SETEXSTATE(&ex, &pltmask, &scalar, &vector, &thresh);

    ex = nextex;
}
}

void
PVSETEXTRACT(int *index, int *exnum, int *pltmsk, int *mask,
             int *ival, float *rvec)
{
    /* called only for programmed extracts */
    if (*index != 101) return;

    /* get the one sub-extract */
    *mask = 1;

    /* set the start location */
    rvec[0] = 0.0; /* X */
    rvec[1] = 0.0; /* Y */
    rvec[2] = 0.0; /* Z */
}

```

7.1.2 Graphics Thread code

```
#include <stdio.h>
#include <stdlib.h>
#include "pV3ser.h"

typedef struct {
    float X; float Y; float Z;} Triad;

void
PVRENDER(float *time)
{
    int    ivec[11], type, exnum, len, cid;
    char   exname[20];
    float  rvec[12];
    void   *ex, *nextex;
    Triad  *streamline;

    /* loop through all defined extracts */
    ex = NULL;
    PV_GETEXTRACT(&ex, &type, &exnum, ivec, rvec, exname, &nextex);
    while (ex != NULL) {
        PV_GETEXTRACT(&ex, &type, &exnum, ivec, rvec, exname, &nextex);

        /* our extract */
        if (type == 101) {
            PV_GETSUB(&ex, 0, 0, (void **) &streamline, &len, &cid);
            /* plot the streamline */
        }

        /* domain surfaces */
        if (type == 5) {
            /* plot domain surfaces */
        }

        ex = nextex;
    }
}
```

7.2 Client Side

It is assumed that most of the **pV3** client-side code has already been constructed. See the **pV3** Programmer's Guide for a complete description of this coupling. The code listed below calculates a streamline using FEL as a programmed extract. It assumes that only one extract has been defined and that extract has only one sub-extract.

It may be necessary to consult the Advanced Programmer's Guide and FEL's documentation to understand this C++ code listing.

```
#include <stdio.h>
#include <stdlib.h>
#include <iostream.h>

#include <FEL.h>
#include "pV3.h"

#define MAX_LENGTH 2000
#define TIMESTEP 0.02

void
pVEXTRACT(int *index, int *exnum, int *reqmask, int *ival, float *rvec)
{
    static int first = 0;
    static FEL_grid *grid;
    static FEL_vector_field *velocity;
    float v[3], *pVel, *buffer, *iblack;
    int i, i3, j, opt, nnodes, nblank;
    int length;
    float streamline[MAX_LENGTH][3];
    vertex_data *pVgrid;
    FEL_bary_pos current_bary;
    FEL_bary_pos last_bary;
    FEL_phys_pos current_position;

    if (first == 0) {
        // instantiate grid object
        grid = new FEL_structured_grid("grid", 1);

        // set the grid type
```

```

grid->set_grid_type(FEL_GRID_P3D_NO_IBLANK);
grid->set_grid_type(FEL_GRID_P3D_SINGLE_ZONE);

// load the grid data
opt = 301;
pV_GETSTRUC(&opt, (void **)&buffer, &nnodes);
pVgrid = (vertex_data *) malloc(nnodes * sizeof(vertex_data));
for (i=0,i3=0; i<nnodes; i++,i3+=3) {
    (pVgrid + i)->x      = buffer[i3  ];
    (pVgrid + i)->y      = buffer[i3+1];
    (pVgrid + i)->z      = buffer[i3+2];
    (pVgrid + i)->iblack = 1;
}
opt = 306;
pV_GETSTRUC(&opt, (void **)&iblack, &nblack);
if (nblack != 0)
    for (i=0; i<nnodes; i++) (pVgrid + i)->iblack=iblack[i];
grid->geom->new_timestep(pVgrid);

// instantiate a vector field
velocity = FEL_make_new_vector_field("velocity", grid, 2);
// select the file type
velocity->set_file_type(FEL_PLOT3D_SOLUTION_FILE);
// load the vector data
opt = 304;
pV_GETSTRUC(&opt, (void **)&pVel, &nnodes);
velocity->new_timestep(pVel);
}
first++;

// initialize the current physical position
length = 0;
current_position.x = rvec[0];
current_position.y = rvec[1];
current_position.z = rvec[2];
current_position.time = 0.0;

// initialize the streamline vertices

```

```

streamline[length][0] = current_position.x;
streamline[length][1] = current_position.y;
streamline[length][2] = current_position.z;

// initialize the barycentric coordinates
grid->phys_to_bary(current_position, current_bary);
last_bary = current_bary;

// Euler streamline computation loop
while (length < MAX_LENGTH-1)
{
    // get the velocity value using last_bary to optimize
    // point location algorithm. Stop if get_value
    // returns 0 as that means we fell off the grid
    if (!velocity->get_value(current_position, last_bary, v)) break;

    // add the velocity times the timestep to the current position
    current_position.x += TIMESTEP * v[0];
    current_position.y += TIMESTEP * v[1];
    current_position.z += TIMESTEP * v[2];
    length++;

    // save the current position as a streamline vertex
    streamline[length][0] = current_position.x;
    streamline[length][1] = current_position.y;
    streamline[length][2] = current_position.z;
}

// send the streamline data to the server
i = 3; j = 0;
length++;
pV_SENDEXR(index, exnum, &j, &i, &length, (float *)streamline);
}

```

A Plotting Masks for Built-in Extracts

The following are additive (or-able) so that the proper attributes can be specified.

A.1 Cut Surfaces

This mask controls the **pV3** scalar tools (planar cuts, programmed cut surfaces, iso-surfaces and domain surfaces) attributes.

- 1 - **Render** - Surface rendering on
- 2 - **Grid** - Mesh display on
- 4 - **Grey** - Surface colored with grey
- 8 - **Threshold** - Surface is thresholded according to the threshold function and limits
- 16 - **Contour** - Contour lines are plotted on the surface
- 32 - **Translucent** - Plot surface using the translucent attribute
- 64 - **Arrows** - Arrow drawing on
- 128 - **Tufts** - Grid of tufts on (dynamic only)
- 256 - **Mapping** - A 2D mapping exists for this surface (domain only)
- 512 - **Probing** - 2D probing is active
- 1024 - **Outline** - Outline drawing is requested (with the mask equal to only this flag)

A.2 StreamLines

This mask controls the **pV3** streamline plotting attributes and therefore the requested sub-extracts.

- 1 - **Render** - StreamLine rendering on
- 2 - **Tube** - Tube rendering on
- 4 - **Grey** - StreamLine drawn with default color
- 8 - **Threshold** - StreamLine is thresholded according to the threshold function and limits
(not currently implemented)
- 16 - **Back** - StreamLine is backward going (can not be active with 32)

- 32 - Fore** - StreamLine goes down stream (can not be active with 16)
- 64 - Ribbon** - Ribbon rendering on (with 2 makes tubes with twist)
- 512 - Particles** - Seeding on
- 2048 - Probing** - StreamLine probe currently active for this StreamLine (Read-only).

A.3 Particles

This mask controls the **pV3** bubble rendering attributes and therefore the requested sub-extracts.

- 1 - Render** - Bubble rendering on
- 2 - Size** - Bubble size based on divergence like tubes - currently not used
- 4 - Grey** - Bubble colored with default color
- 16 - Time** - Bubbles are colored with the time of spawning
- 32 - Time Lines** - Plot lines between particles in the same group

A.4 Vector Clouds

- 1 - Render** - Vector cloud rendering on
- 4 - Grey** - Vector cloud colored with default color